Received: 20 July 2023 | Revised: 30 August 2023 | Accepted: 31 August 2023

DOI: 10.1002/psp4.13047

PERSPECTIVE

Quantitative systems pharmacology in the age of artificial intelligence

Benjamin Ribba

Roche Pharma Research and Early Development (pRED), Roche Innovation Center, F. Hoffmann-La Roche Ltd., Basel, Switzerland

Correspondence
Benjamin Ribba, Roche Pharma Research and Early Development (pRED), Roche Innovation Center, F. Hoffmann-La Roche Ltd., Basel, Switzerland.
Email: benjamin.ribba@roche.com

In this perspective, we briefly review the state-of-the-art covering the interface between quantitative systems pharmacology (QSP) and artificial intelligence (AI) and machine learning (ML); in particular, how AI/ML are conceived and used as many attempts to address methodological pain points of QSP. In a second part, we invite the reader to step out from this discipline-centric view and discuss a paradigm shift consisting of repurposing AI into QSP.

In 2022, the *Journal of Pharmacokinetic and Pharmacodynamic* published a special issue edited by Cho, Zhang, and Bonate with 10 scientific articles illustrating ways for coupling quantitative systems pharmacology (QSP) and artificial intelligence (AI). The review from Zang et al. summarizes the current state-of-the-art[1] and explains that AI/machine learning (ML) are currently used for four main applications related to QSP: parameter estimation, model structure, complexity reduction, and virtual population generations (see Figure 1).

What may strike the reader is that – although there are good reasons to think that improving on such technical domains will increase the impact of QSP – these four domains speak toward technical dimensions of QSP as a discipline. Estimation of parameters, writing down optimal model structures, reducing model complexity, or properly generating virtual population are terms which will talk to modeling experts but probably only to them.

Parameter estimation is known to be an issue for QSP because the models are often too large with respect to available data. QSP modelers are typically working as engineers; writing down equations to describe complex systems with a high level of granularity, although the elementary biological processes are unknown or uncertain. In addition to that, the clinical end points we aim to predict are often remote from the biological level we expend effort to accurately describe. We, as a community, are trying to address these issues with AI/ML. On one hand, big data, such as omics and/or imaging data, can be treated as inputs of supervised learning and QSP model parameters as outputs (see ref. 2 as an explanation of such an approach in oncology). On the other hand, QSP model parameters could be linked to clinical endpoints with a similar formalism where the QSP model parameters are the inputs and the clinical end points are the outputs to be predicted.

Finding the optimal structural model is also known as an issue, as often, the model structure is designed by hand; and thus subjected to subjectivity coming from whomever wrote it. We can use data-driven approaches to inform model structure through Boolean networks, Bayesian networks, or even physics-based AI. The approach has received attention especially for safety prediction, as shown by a recent initiative to identify the molecular players involved in drug-induced peripheral neuropathy.[3]

Regarding model complexity, authors propose to apply feature selection techniques to reduce a priori QSP model complexity to lead to a parsimonious model more amenable to QSP tasks (see ref. 1 for further information). Last but not least, the creation of virtual populations (VPs) has been rightly the focus of much interest and approaches from ML, such as prevalence weighting and Markov Chain Monte

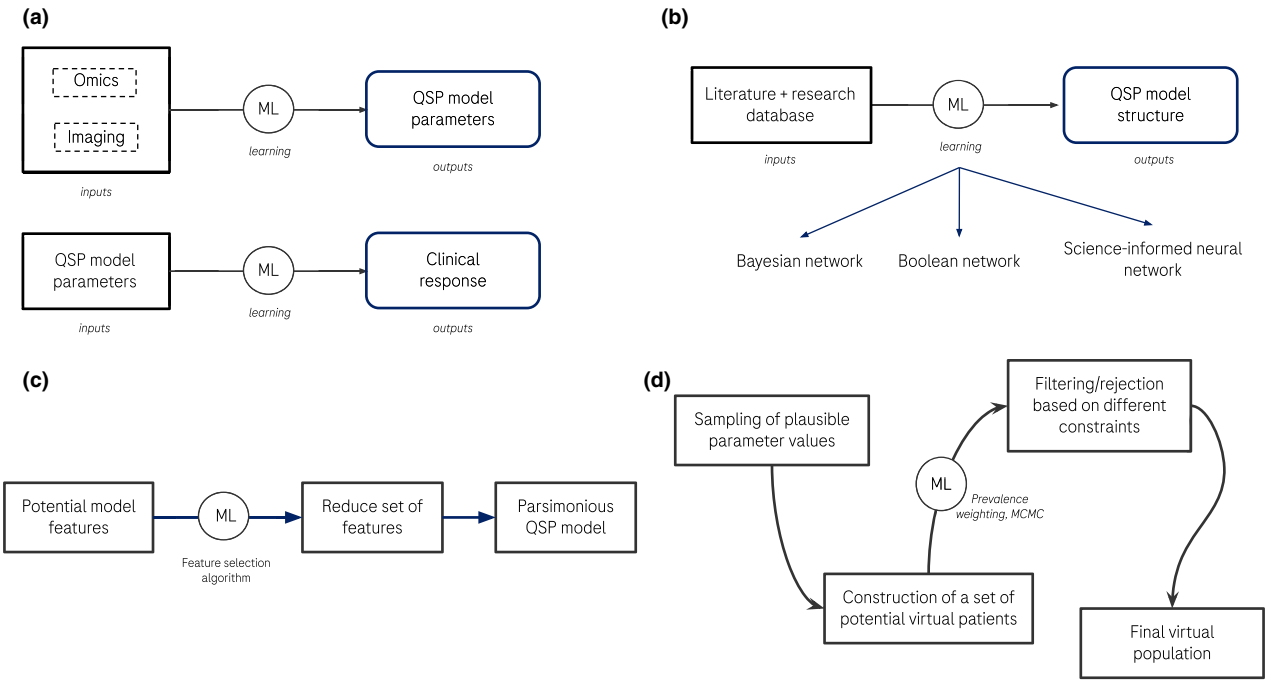

FIGURE 1 Schematic representation of the current use of AI/ML in QSP. (**a**) Is for the parameter estimation; (**b**) is for the model structure creation; (**c**) is for the dimension reduction; (**d**) is for the virtual population creations. AI, artificial intelligence; MCMC, Markov Chain Monte Carlo; ML, machine learning; QSP, quantitative systems pharmacology.

Carlo, are proposed to identify relevant parameters from a plausible ensemble in order to generate virtual patients resembling clinical cohorts.[4] In addition to that, AI approaches are proposed to speed up simulations of computational-demanding QSP models for efficient VP generation.[5]

This list of proposed added values is very much "QSP-centric." The paradigm is centered within QSP and we ask ourselves how we can take advantage of AI to solve some of the most important methodological problems inherent to QSP. There is nothing wrong in this and the reported issues need to find methodological solutions, however, it should not be the sink of all our attention and should not deviate our community from two fundamental things:

- The focus on QSP's mission, which is to advance the development and safe use of medicine;
- The focus on QSP's approach or "philosophy" (rather than on the "tools" to do QSP). The "philosophy" is to be quantitative, bring a system view, and apply interpretable pharmacological principles of the action/effect of therapeutic interventions.

Regarding the key pain points of drug development and safe use of medicine, two of them immediately come to mind.

First, the multiscale nature of clinical end points, which in some areas incorporate aspects of daily living, such as the clinical dementia rating in Alzheimer's disease (AD) or the unified Parkinsons's disease rating scale (UPDRS). Leveraging such data through modeling approaches is notoriously challenging. This difficulty will probably spread to other therapeutic areas given the increasing availability of patients' reported outcomes.

We sometimes highlight the problem of these end points being "subjective." In the author's opinion, the subjectivity is less of a problem than the multiscale and multidimensional nature of such measures. We are lacking the methods to bridge to the level of these end points from the traditional biological level, which is typically the focus of our models. This means that we have an urgency to double down on innovation in disease progression modeling which can in some way be descriptive of the level of information contained in the clinical end points we aim to predict.

Second, we lack validated biomarkers, and we have a responsibility to look beyond the classical markers measured in blood or tissue to potential surrogates of drug action. Still, in the field of neuroscience, these would be markers of the cognitive, behavioral, and functional levels.

How do we, as a community, propose to address these pain points?

The good news is that AI can be repurposed/reframed into QSP to stick to these two fundamental points of focus. Indeed, several recently published reports suggest taking advantage of AI algorithm capability to perform well on human tasks as a basis for "mechanistic" disease models.

An illustrative example is around the reframing of AI algorithms performing well at facial recognition and

their analysis to shed new light on mechanisms happening in patients suffering from AD and losing facial recognition capabilities.[6] After having trained a model to perform well (mimicking a healthy individual) the authors perturbed the model in two different ways, either modifying the weight to mimic brain signal weakening or cutting on some neural network nodes to mimic brain lesions. Simulations are performed with this perturbed model and the results of these simulations were put into perspective with biomarker data from the AD neuroimaging initiative. This paper presents the proposed shift of paradigm very clearly. However, because such deep neural network models are lacking interpretability, the question whether indeed the learning coming from "reverse engineering" the AI model could be informative on the actual pathological processes happening in the human brain remains fully open.

Following a similar strategy with an AI technique called reinforcement learning (RL) could even be more promising because the technique is supported by theoretical arguments. Indeed, it has been shown experimentally that RL can be a valid model for analyzing some aspects of humans and animals' learning.

It is with the following words that Richard Sutton and Andrew Barto start their book on RL[7]:

> When an infant learns how to walk, fall and try again, it has no explicit teacher but it does have a direct sensorimotor connection to its environment. Exercising this connection produces a wealth of information about cause and effect, about consequences of actions, and about what to do to achieve goals.

RL mimics the process of learning and consists of knowing what to do – how to map situations to actions – so as to maximize a numerical reward signal. RL is often represented with Markov Decision Process formalism which consists of an agent that takes actions, can sense its environment, and has clarity about a goal to achieve. The process can be "computerized" which has led to many applications including for health. Computationally, RL aims at estimating a value function iteratively, through experience, which is updated by a term which is called "temporal difference" (TD). The TD is the difference between the currrent expectation of a value of a state and what the value is, after experience. If TD is positive, the value of the state can be increased, if TD is negative, it should be then decreased. This process can happen until convergence.

RL has been highlighted as an interesting approach for precision dosing of pharmacology or non-pharmacological interventions and for computational psychiatry.[8] Before reporting a definition of computational psychiatry, let us illustrate it with two examples.

In a paper published in 2004 in *Science*, David Redish repurposed RL to simulate addiction.[9] Following the previous explanation on RL and the concept of TD, addiction is modeled by a TD which is always positive when some specific actions are taken (i.e., the one the subject is addicted to). By modifying the RL algorithm in such a way, actions for which the subject is addicted to will always result in positive TD, which will increase the corresponding value. In turn, such actions are most likely to be chosen by the subject.

Earlier, we discussed the need for validated biomarkers and RL have been shown to be very useful to analyze data from cognitive testing (see ref. 8 for further examples).

Computational psychiatry is a nascent scientific area defined as a way to characterize mental dysfunction in terms of aberrant brain computations. In the author's opinion, computational psychiatry, when coupled with (pharmacological) description of intervention, is nothing else than a QSP approach. To how QSP is envisioned today in our scientific community, RL reinforces the "system" component and could be highly beneficial in therapeutic areas, such as neurodegeneration and psychiatry for which QSP has limited impact today for the reasons discussed earlier.

In conclusion, we discussed two ways to see the topic of QSP interfacing with AI. The first is a very much inner view, QSP-centric, where we leverage AI to do QSP. What seems as an underdeveloped area which can have high potential is to step out of this box and to focus on the key points and reframe AI into QSP (Figure 2).

Embracing such a paradigm could make QSP even more critical to advance the development of new therapeutic

FIGURE 2 Left: The inner view: QSP-discipline centric. Right: The outer view: repurposing AI into QSP. AI, artificial intelligence; QSP, quantitative systems pharmacology.

modalities. QSP is doing a tremendous role in supporting the development and safe use of molecular therapy. Recently, a new therapeutic modality is emerging, digital therapeutics (DTx), and known as software applications to treat, manage, or prevent a condition. DTx are considered as very promising modalities in areas of behaviors, mental disorders, sleep disorders, and pain. We have recently reviewed the field and identified that application of clinical pharmacology and modeling principles would greatly benefit the development of DTx, in particular, for better characterizing the mechanism of action, optimize intervention content, and identify the right dose and the right patients.[10]

QSP could drive the development of such modalities, becoming instrumental in overcoming another identified limitation of current therapeutic innovation in this field, namely the high selectivity of drug candidates not reflecting the complex interactions of different brain circuits.

So, are we up to the challenge?

ACKNOWLEDGMENTS

The author acknowledges Dr. Cristina Santini and the UK QSP network for their invitation to present on the topic at the 4th QSP UK Exchange Workshop on July 12, 2023; and the workshop participants too for their valuable questions and feedback which helped in writing the present perspective. The author thanks Dr. Lucy Hutchinson for the scientific exchanges and help in finalizing the manuscript.

FUNDING INFORMATION

No funding was received for this work.

CONFLICT OF INTEREST STATEMENT

The author is an employee and shareholder of F. Hoffmann-La Roche Ltd.

REFERENCES

1. Zhang T, Androulakis IP, Bonate P, et al. Two heads are better than one: current landscape of integrating QSP and machine learning: an ISoP QSP SIG white paper by the working group on the integration of quantitative systems pharmacology and machine learning. *J Pharmacokinet Pharmacodyn*. 2022;49(1):5-18.
2. Benzekry S. Artificial intelligence and mechanistic modeling for clinical decision making in oncology. *Clin Pharmacol Ther*. 2020;108(3):471-486.
3. Hur J, Guo AY, Loh WY, Feldman EL, Bai JP. Integrated systems pharmacology analysis of clinical drug-induced peripheral neuropathy. *CPT Pharmacometrics Syst Pharmacol*. 2014;3(5):e114.
4. Rieger TR, Allen RJ, Musante CJ. Modeling is data driven: use it for successful virtual patient generation. *CPT Pharmacometrics Syst Pharmacol*. 2021;10(5):393-394.
5. Myers RC, Augustin F, Huard J, Friedrich CM. Using Machine Learning Surrogate Modeling for faster QSP VP Cohort Generation. *CPT Pharmacometrics Syst Pharmacol*. 2023;12(8):1047-1059.
6. Singh G, Ramanathan M. Repurposing artificial intelligence tools for disease modeling: case study of face recognition deficits in neurodegenerative diseases. *Clin Pharmacol Ther*. 2023;114(4):862-873.
7. Sutton R, Barto A. *Reinforcement Learning: an Introduction*. Second ed. MIT Press; 2018.
8. Ribba B. Reinforcement learning as an innovative model-based approach: examples from precision dosing, digital health and computational psychiatry. *Front Pharmacol*. 2022;13:1094281.
9. Redish AD. Addiction as a computational process gone awry. *Science*. 2004;306(5703):1944-1947.
10. Ribba B, Peck R, Hutchinson L, Bousnina I, Motti D. Digital therapeutics as a new therapeutic modality: a review from the perspective of clinical pharmacology. *Clin Pharmacol Ther*. 2023;114:578-590.

How to cite this article: Ribba B. Quantitative systems pharmacology in the age of artificial intelligence. *CPT Pharmacometrics Syst Pharmacol*. 2023;12:1823-1826. doi:10.1002/psp4.13047